arXiv:1510.08512v1 [stat.ML] 28 Oct 2015

Robust Gaussian Graphical Modeling with the Trimmed Graphical Lasso

Eunho Yang and Aurélie C. Lozano
IBM T.J. Watson Research Center

Abstract

Gaussian Graphical Models (GGMs) are popular tools for studying network structures. However, many modern applications such as gene network discovery and social interactions analysis often involve high-dimensional noisy data with outliers or heavier tails than the Gaussian distribution. In this paper, we propose the Trimmed Graphical Lasso for robust estimation of sparse GGMs. Our method guards against outliers by an implicit trimming mechanism akin to the popular Least Trimmed Squares method used for linear regression. We provide a rigorous statistical analysis of our estimator in the high-dimensional setting. In contrast, existing approaches for robust sparse GGMs estimation lack statistical guarantees. Our theoretical results are complemented by experiments on simulated and real gene expression data which further demonstrate the value of our approach.

1 Introduction

Gaussian graphical models (GGMs) form a powerful class of statistical models for representing distributions over a set of variables [1]. These models employ undirected graphs to encode conditional independence assumptions among the variables, which is particularly convenient for exploring network structures. GGMs are widely used in variety of domains, including computational biology [2], natural language processing [3], image processing [4, 5, 6], statistical physics [7], and spatial statistics [8]. In many modern applications, the number of variables p can exceed the number of observations n. For instance, the number of genes in microarray data is typically larger than the sample size. In such high-dimensional settings, sparsity constraints are particularly pertinent for estimating GGMs, as they encourage only a few parameters to be non-zero and induce graphs with few edges. The most widely used estimator among others (see e.g. [9]) minimizes the Gaussian negative log-likelihood regularized by the ℓ_1 norm of the entries (or the off-diagonal entries) of the precision matrix (see [10, 11, 12]). This estimator enjoys strong statistical guarantees (see e.g. [13]). The corresponding optimization problem is a log-determinant program that can be solved with interior point methods [14] or by co-ordinate descent algorithms [11, 12]. Alternatively neighborhood selection [15, 16] can be employed to estimate conditional independence relationships separately for each node in the graph, via Lasso linear regression, [17]. Under certain assumptions, the sparse GGM structure can still be recovered even under high-dimensional settings.

The aforementioned approaches rest on a fundamental assumption: the multivariate normality of the observations. However, outliers and corruption are frequently encountered in high-dimensional data (see e.g. [18] for gene expression data). Contamination of a few observations can drastically affect the quality of model estimation. It is therefore imperative to devise procedures

that can cope with observations deviating from the model assumption. Despite this fact, little attention has been paid to robust estimation of high-dimensional graphical models. Relevant work includes [19], which leverages multivariate t-distributions for robustified inference and the EM algorithm. They also propose an alternative t-model which adds flexibility to the classical t but requires the use of Monte Carlo EM or variational approximation as the likelihood function is not available explicitly. Another pertinent work is that of [20] which introduces a robustified likelihood function. A two-stage procedure is proposed for model estimation, where the graphical structure is first obtained via coordinate gradient descent and the concentration matrix coefficients are subsequently re-estimated using iterative proportional fitting so as to guarantee positive definiteness of the final estimate.

In this paper, we propose the *Trimmed Graphical Lasso* method for robust Gaussian graphical modeling in the sparse high-dimensional setting. Our approach is inspired by the classical Least Trimmed Squares method used for robust linear regression [21], in the sense that it disregards the observations that are judged less reliable. More specifically the Trimmed Graphical Lasso seeks to minimize a weighted version of the negative log-likelihood regularized by the ℓ_1 penalty on the concentration matrix for the GGM and under some simple constraints on the weights. These weights implicitly induce the trimming of certain observations. Our key contributions can be summarized as follows.

- We introduce the Trimmed Graphical Lasso formulation, along with two strategies for solving the objective. One involves solving a series of graphical lasso problems; the other is more efficient and leverages composite gradient descent in conjunction with partial optimization.
- As our key theoretical contribution, we provide statistical guarantees on the consistency of our estimator. To the best of our knowledge, this is in stark contrast with prior work on robust sparse GGM estimation (e.g. [19, 20]) which do not provide any statistical analysis.
- Experimental results under various data corruption scenarios further demonstrate the value of our approach.

2 Problem Setup and Robust Gaussian Graphical Models

Notation. For matrices $U \in \mathbb{R}^{p\times p}$ and $V \in \mathbb{R}^{p\times p}$, $\langle\langle U, V\rangle\rangle$ denotes the trace inner product $\mathrm{tr}(A\,B^T)$. For a matrix $U \in \mathbb{R}^{p\times p}$ and parameter $a \in [1,\infty]$, $\|U\|_a$ denotes the element-wise ℓ_a norm, and $\|U\|_{a,\mathrm{off}}$ does the element-wise ℓ_a norm only for off-diagonal entries. For example, $\|U\|_{1,\mathrm{off}} := \sum_{i\neq j}|U_{ij}|$. Finally, we use $\|U\|_{\mathrm{F}}$ and $|\!|\!|U|\!|\!|_2$ to denote the Frobenius and spectral norms, respectively.

Setup. Let $X = (X_1, X_2, \ldots, X_p)$ be a zero-mean Gaussian random field parameterized by $p \times p$ concentration matrix Θ^*:

$$\mathbb{P}(X;\Theta^*) = \exp\Big(-\frac{1}{2}\langle\langle \Theta^*, XX^\top\rangle\rangle - A(\Theta^*)\Big) \tag{1}$$

where $A(\Theta^*)$ is the log-partition function of Gaussian random field. Here, the probability density function in (1) is associated with p-variate Gaussian distribution, $N(0,\Sigma^*)$ where $\Sigma^* = (\Theta^*)^{-1}$.

Given n i.i.d. samples $\{X^{(1)}, \ldots, X^{(n)}\}$ from high-dimensional Gaussian random field (1), the standard way to estimate the inverse covariance matrix is to solve the ℓ_1 regularized maximum likelihood estimator (MLE) that can be written as the following regularized *log-determinant* program:

$$\underset{\Theta \in \Omega}{\text{minimize}} \left\langle\!\!\left\langle \Theta, \frac{1}{n}\sum_{i=1}^{n} X^{(i)}(X^{(i)})^\top \right\rangle\!\!\right\rangle - \log\det(\Theta) + \lambda\|\Theta\|_{1,\text{off}} \tag{2}$$

where Ω is the space of the symmetric positive definite matrices, and λ is a regularization parameter that encourages a *sparse* graph model structure.

In this paper, we consider the case where the number of random variables p may be substantially larger than the number of sample size n, however, the concentration parameter of the underlying distribution is sparse:

(C-1) The number of non-zero off-diagonal entries of Θ^* is at most k, that is $|\{\Theta^*_{ij} : \Theta^*_{ij} \neq 0 \text{ for } i \neq j\}| \leq k$.

Now, suppose that n samples are drawn from this underlying distribution (1) with true parameter Θ^*. We further allow some samples are corrupted and not drawn from (1). Specifically, the set of sample indices $\{1, 2, \ldots, n\}$ is separated into two disjoint subsets: if i-th sample is in the set of "good" samples, which we name G, then it is a genuine sample from (1) with the parameter Θ^*. On the other hand, if the i-th sample is in the set of "bad" samples, B, the sample is corrupted. The identifications of G and B are hidden to us. However, we naturally assume that only a small number of samples are corrupted:

(C-2) Let h be the number of good samples: $h := |G|$ and hence $|B| = n - h$. Then, we assume that larger portion of samples are genuine and uncorrupted so that $\frac{|G|-|B|}{|G|} \geq \alpha$ where $0 < \alpha \leq 1$. If we assume that 40% of samples are corrupted, then $\alpha = \frac{0.6n-0.4n}{0.6n} = \frac{1}{3}$.

In later sections, we will derive a *robust* estimator for corrupted samples of sparse Gaussian graphical models and provide statistical guarantees of our estimator under the conditions (C-1) and (C-2).

2.1 Trimmed Graphical Lasso

We now propose a *Trimmed Graphical Lasso* for robust estimation of sparse GGMs:

$$\begin{aligned} \underset{\Theta \in \Omega, w}{\text{minimize}} \; & \left\langle\!\!\left\langle \Theta, \frac{1}{h}\sum_{i=1}^{n} w_i X^{(i)}(X^{(i)})^\top \right\rangle\!\!\right\rangle - \log\det(\Theta) + \lambda\|\Theta\|_{1,\text{off}} \\ \text{s.t. } & w \in [0,1]^n \,,\, \mathbf{1}^\top w = h \,,\, \|\Theta\|_1 \leq R \end{aligned} \tag{3}$$

where λ is a regularization parameter to decide the sparsity of our estimation, and h is another parameter, which decides the number of samples (or sum of weights) used in the training. h is ideally set as the number of uncorrupted samples in G, but practically we can tune the parameter h by cross-validation. Here, the constraint $\|\Theta\|_1 \leq R$ is required to analyze this non-convex optimization problem as discussed in [22]. For another tuning parameter R, any positive real value would be sufficient for R as long as $\|\Theta^*\|_1 \leq R$. Finally note that when h is fixed as n (and R is set as infinity), the optimization problem (3) will be simply reduced to the vanilla ℓ_1 regularized MLE for sparse GGM without concerning outliers.

The optimization problem (3) is convex in w as well as in Θ, however this is not the case jointly. Nevertheless, we will show later that **any** local optimum of (3) is guaranteed to be strongly consistent under some fairly mild conditions.

Algorithm 1 Trimmed Graphical Lasso in (3)

Initialize $\Theta^{(0)}$ (e.g. $\Theta^{(0)} = (S + \lambda I)^{-1}$)
repeat
Compute $w^{(t)}$ given $\Theta^{(t-1)}$, by assigning a weight of one to the h observations with lowest negative log-likelihood and a weight of zero to the remaining ones.
$\nabla \mathcal{L}^{(t)} \leftarrow \frac{1}{h} \sum_{i=1}^{n} w_i^{(t)} X^{(i)} (X^{(i)})^\top - (\Theta^{(t-1)})^{-1}$
Line search. Choose $\eta^{(t)}$ (See Nesterov (2007) for a discussion of how the stepsize may be chosen), checking that the following update maintains positive definiteness. This can be verified via Cholesky factorization (as in [23]).
Update. $\Theta^{(t)} \leftarrow \mathcal{S}_{\eta^{(t)}\lambda}(\Theta^{(t-1)} - \eta^{(t)} \nabla \mathcal{L}^{(t)})$, where $\mathcal{S}$ is the soft-thresholding operator: $[S_\nu(U)]_{i,j} = \text{sign}(U_{i,j}) \max(|U_{i,j}| - \nu, 0)$ and is only applied to the off-diagonal elements of matrix U.
Compute $(\Theta^{(t)})^{-1}$ reusing the Cholesky factor.
until stopping criterion is satisfied

Optimization As we briefly discussed above, the problem (3) is not jointly convex but biconvex. One possible approach to solve the objective of (3) thus is to alternate between solving for Θ with fixed w and solving for w with fixed Θ. Given Θ, solving for w is straightforward and boils down to assigning a weight of one to the h observations with lowest negative log-likelihood and a weight of zero to the remaining ones. Given w, solving for Θ can be accomplished by any algorithm solving the "vanilla" graphical Lasso program, e.g. [11, 12]. Each step solves a convex problem hence the objective is guaranteed to decrease at each iteration and will converge to a local minima.

A more efficient optimization approach can be obtained by adopting a partial minimization strategy for Θ: rather than solving to completion for Θ each time w is updated, one performs a single step update. This approach stems from considering the following equivalent reformulation of our objective:

$$\begin{aligned} &\underset{\Theta \in \Omega}{\text{minimize}} \left\langle\!\!\left\langle \Theta, \frac{1}{h} \sum_{i=1}^{n} w_i(\Theta) X^{(i)} (X^{(i)})^\top \right\rangle\!\!\right\rangle - \log\det(\Theta) + \lambda \|\Theta\|_{1,\text{off}} \\ &\text{s.t. } w_i(\Theta) = \underset{w \in [0,1]^n,\ \mathbf{1}^\top w = h}{\text{argmin}} \left\langle\!\!\left\langle \Theta, \frac{1}{h} \sum_{i=1}^{n} w_i X^{(i)} (X^{(i)})^\top \right\rangle\!\!\right\rangle, \quad \|\Theta\|_1 \le R, \end{aligned} \tag{4}$$

On can then leverage standard first-order methods such as projected and composite gradient descent [24] that will converge to local optima. The overall procedure is depicted in Algorithm 1. Therein we assume that we pick R sufficiently large, so one does not need to enforce the constraint $\|\Theta\|_1 \le R$ explicitly. If needed the constraint can be enforced by an additional projection step [22].

3 Statistical Guarantees of Trimmed Graphical Lasso

One of the main contributions of this paper is to provide the statistical guarantees of our Trimmed Graphical Lasso estimator for GGMs. The optimization problem (3) is non-convex, and therefore the gradient-type methods solving (3) will find estimators by local minima. Hence, our theory in this section provides the statistical error bounds on **any** local minimum measured by $\|\cdot\|_{\text{F}}$ and $\|\cdot\|_{1,\text{off}}$ norms simultaneously.

Suppose that we have some local optimum $(\widetilde{\Theta}, \widetilde{w})$ of (3) by arbitrary gradient-based method. While Θ^* is fixed unconditionally, we define w^* as follows: for a sample index $i \in G$, w_i^* is simply

set to $\widetilde{w}_i$ so that $w_i^* - \widetilde{w}_i = 0$. Otherwise for a sample index $i \in B$, we set $w_i^* = 0$. Hence, w^* is dependent on $\widetilde{w}$.

In order to derive the upper bound on the Frobenius norm error, we first need to assume the standard restricted strong convexity condition of (3) with respective to the parameter Θ:

(C-3) (Restricted strong convexity condition) Let Δ be an arbitrary error of parameter Θ. That is, $\Delta := \Theta - \Theta^*$. Then, for any possible error Δ such that $\|\Delta\|_{\mathrm{F}} \leq 1$,

$$\left\langle\!\!\left\langle (\Theta^*)^{-1} - (\Theta^* + \Delta)^{-1}, \Delta \right\rangle\!\!\right\rangle \geq \kappa_l \|\Delta\|_{\mathrm{F}}^2 \tag{5}$$

where κ_l is a curvature parameter.

Note that in order to guarantee the Frobenius norm-based error bounds, (C-3) is required even for the vanilla Gaussian graphical models *without* outliers, which has been well studied by several works such as the following lemma:

Lemma 1 (Section B.4 of [22]). *For any $\Delta \in \mathbb{R}^{p \times p}$ such that $\|\Delta\|_{\mathrm{F}} \leq 1$,*

$$\left\langle\!\!\left\langle (\Theta^*)^{-1} - (\Theta^* + \Delta)^{-1}, \Delta \right\rangle\!\!\right\rangle \geq \left(|\!|\!|\Theta^*|\!|\!|_2 + 1\right)^{-2} \|\Delta\|_{\mathrm{F}}^2 ,$$

thus (C-3) holds with $\kappa_l = \left(|\!|\!|\Theta^*|\!|\!|_2 + 1\right)^{-2}$.

While (C-3) is a standard condition that is also imposed for the conventional estimators under clean set of of samples, we additionally require the following condition for a successful estimation of (3) on corrupted samples:

(C-4) Consider arbitrary local optimum $(\widetilde{\Theta}, \widetilde{w})$. Let $\widetilde{\Delta} := \widetilde{\Theta} - \Theta^*$ and $\widetilde{\Gamma} := \widetilde{w} - w^*$. Then,

$$\left| \left\langle\!\!\left\langle \frac{1}{h} \sum_{i=1}^{n} \widetilde{\Gamma}_i X^{(i)} (X^{(i)})^\top, \widetilde{\Delta} \right\rangle\!\!\right\rangle \right| \leq \tau_1(n,p) \|\widetilde{\Delta}\|_{\mathrm{F}} + \tau_2(n,p) \|\widetilde{\Delta}\|_1$$

with some positive quantities $\tau_1(n,p)$ and $\tau_2(n,p)$ on n and p. These will be specified below for some concrete examples.

(C-4) can be understood as a structural incoherence condition between the model parameter Θ and the weight parameter w. Such a condition is usually imposed when analyzing estimators with multiple parameters (for example, see [25] for a robust linear regression estimator). Since w^* is defined depending on $\widetilde{w}$, each local optimum has its own (C-4) condition. We will see in the sequel that under some reasonable cases, this condition for any local optimum holds with high probability. Also note that for the case with clean samples, the condition (C-4) is trivially satisfied since $\widetilde{\Gamma}_i = 0$ for all $i \in \{1, \ldots, n\}$ and hence the LHS becomes 0.

Armed with these conditions, we now state our main theorem on the error bounds of our estimator (3):

Theorem 1. *Consider corrupted Gaussian graphical models. Let $(\widetilde{\Theta}, \widetilde{w})$ be an* **any** *local optimum of M-estimator (3). Suppose that $(\widetilde{\Theta}, \widetilde{w})$ satisfies the condition (C-4). Suppose also that the regularization parameter λ in (3) is set such that*

$$4 \max \left\{ \left\| \frac{1}{h} \sum_{i=1}^{n} w_i^* X^{(i)} (X^{(i)})^\top - (\Theta^*)^{-1} \right\|_\infty, \tau_2(n,p) \right\} \leq \lambda \leq \frac{\kappa_l - \tau_1(n,p)}{3R}. \tag{6}$$

Then, this local optimum $(\widetilde{\Theta}, \widetilde{w})$ *is guaranteed to be consistent as follows:*

$$\|\widetilde{\Theta} - \Theta^*\|_{\mathrm{F}} \leq \frac{1}{\kappa_l}\Big(\frac{3\lambda\sqrt{k+p}}{2} + \tau_1(n,p)\Big) \quad \textit{and}$$
$$\|\widetilde{\Theta} - \Theta^*\|_{1,\mathrm{off}} \leq \frac{2}{\lambda\,\kappa_l}\Big(3\lambda\sqrt{k+p} + \tau_1(n,p)\Big)^2 . \tag{7}$$

The statement in Theorem 1 holds deterministically, and the probabilistic statement comes where we show (C-4) and (6) for a given $(\widetilde{\Theta}, \widetilde{w})$ are satisfied. Note that, defining $\mathcal{L}(\Theta, w) := \big\langle\!\big\langle \Theta,\, \frac{1}{h}\sum_{i=1}^n w_i X^{(i)}(X^{(i)})^\top \big\rangle\!\big\rangle - \log\det(\Theta)$, it is a standard way of choosing λ based on $\|\nabla_\Theta \mathcal{L}\big(\Theta^*, w^*\big)\|_\infty$ (see [26], for details). Also it is important to note that the term $\sqrt{k+p}$ captures the relation between element-wise ℓ_1 norm and the error norm $\|\cdot\|_{\mathrm{F}}$ including *diagonal entries*. Due to the space limit, the proof of Theorem 1 (and all other proofs) are provided in the Supplements.

Now, it is natural to ask how easily we can satisfy the conditions in Theorem 1. Intuitively it is impossible to recover true parameter by weighting approach as in (3) when the amount of corruptions exceeds that of normal observation errors.

To this end, suppose that we have some upper bound on the corruptions:

(C-B1) For some function $f(\cdot)$, we have $\big(\|\!|X^B|\!\|_2\big)^2 \leq f(X^B)\sqrt{h\log p}$

where X^B denotes the sub-design matrix in $\mathbb{R}^{|B|\times p}$ corresponding to outliers. Under this assumption, we can properly choose the regularization parameter λ satisfying (6) as follows:

Corollary 1. *Consider corrupted Gaussian graphical models with conditions (C-2) and (C-B1). Suppose that we choose the regularization parameter*

$$\lambda = 4\max\left\{8(\max_i \Sigma^*_{ii})\sqrt{\frac{10\tau\log p}{h-|B|}} + \frac{|B|}{h}\|\Sigma^*\|_\infty\,,\; f(X^B)\sqrt{\frac{\log p}{h}}\right\} \leq \frac{\kappa_l - f(X^B)\sqrt{\frac{|B|\log p}{h}}}{3R}\,.$$

Then, **any** *local optimum of (3) is guaranteed to satisfy (C-4) and have the error bounds in (7) with probability at least* $1 - c_1\exp(-c_1' h\lambda^2)$ *for some universal positive constants* c_1 *and* c_1'.

If we further assume the number of corrupted samples scales with $\sqrt{n}$ at most :

(C-B2) $|B| \leq a\sqrt{n}$ for some constant $a \geq 0$,

then we can derive the following result as another corollary of Theorem 1:

Corollary 2. *Consider corrupted Gaussian graphical models. Suppose that the conditions (C-2), (C-B1) and (C-B2) hold. Also suppose that the regularization parameter* λ *is set as* $c\sqrt{\frac{\log p}{n}}$ *where* $c := 4\max\big\{16(\max_i \Sigma^*_{ii})\sqrt{5\tau} + \frac{2a\|\Sigma^*\|_\infty}{\sqrt{\log p}}\,,\; \sqrt{2}f(X^B)\big\}$. *Then, if the sample size* n *is lower bounded as*

$$n \geq \max\left\{16a^2\,,\; \big(\|\!|\Theta^*|\!\|_2 + 1\big)^4\Big(3Rc + f(X^B)\sqrt{2|B|}\Big)^2(\log p)\right\},$$

then **any** *local optimum of (3) is guaranteed to satisfy (C-4) and have the following error bound:*

$$\|\widetilde{\Theta} - \Theta^*\|_{\mathrm{F}} \leq \frac{1}{\kappa_l}\left(\frac{3c}{2}\sqrt{\frac{(k+p)\log p}{n}} + f(X^B)\sqrt{\frac{2|B|\log p}{n}}\right) \tag{8}$$

with probability at least $1 - c_1\exp(-c_1' h\lambda^2)$ *for some universal positive constants* c_1 *and* c_1'.

Note that the $\|\cdot\|_{1,\text{off}}$ norm based error bound also can be easily derived using the selection of λ from (7). Corollary 2 reveals an interesting result: even when $O(\sqrt{n})$ samples out of total n samples are corrupted, our estimator (3) can successfully recover the true parameter with guaranteed error in (8). The first term in this bound is $O\left(\sqrt{\frac{(k+p)\log p}{n}}\right)$ which exactly recovers the Frobenius error bound for the case without outliers (see [13, 22] for example). Due to the outliers, we have the performance degrade with the second term, which is $O\left(\sqrt{\frac{|B|\log p}{n}}\right)$. To the best of our knowledge, this is the first statistical error bounds on the parameter estimation for Gaussian graphical models with outliers. Also note that Corollary 1 only concerns on any local optimal point derived by an arbitrary optimization algorithm. For the guarantees of multiple local optima simultaneously, we may use a union bound from the corollary.

When Outliers Follow a Gaussian Graphical Model Now let us provide a concrete example and show how $f(X^B)$ in (C-B1) is precisely specified in this case:

(C-B3) Outliers in the set B are drawn from another Gaussian graphical model (1) with a parameter $(\Sigma_B)^{-1}$.

This can be understood as the Gaussian mixture model where the most of the samples are drawn from $(\Theta^*)^{-1}$ that we want to estimate, and small portion of samples are drawn from Σ_B. In this case, Corollary 2 can be further shaped as follows:

Corollary 3. *Suppose that the conditions (C-2), (C-B2) and (C-B3) hold. Then the statement in Corollary 2 holds with* $f(X^B) := \frac{4\sqrt{2}a\left(1+\sqrt{\log p}\right)^2 \|\Sigma_B\|_2}{\sqrt{\log p}}$.

4 Experiments

In this section we corroborate the performance of our Trimmed Graphical Lasso (*trim-glasso*) algorithm on simulated data. We compare against *glasso*: the vanilla Graphical Lasso [11]; the *t-Lasso* and *t*-lasso* methods [19], and *robust-LL*: the robustified-likelihood approach of [20].

4.1 Simulated data

Our simulation setup is similar to [20] and is a akin to gene regulatory networks. Namely we consider four different scenarios where the outliers are generated from models with different graphical structures. Specifically, each sample is generated from the following mixture distribution:

$$y_k \sim (1-p_0)N_p(0,\Theta^{-1}) + \frac{p_0}{2}N_p(-\mu,\Theta_o^{-1}) + \frac{p_0}{2}N_p(\mu,\Theta_o^{-1}), \quad k=1,\ldots,n,$$

where $p_o = 0.1, n = 100$, and $p = 150$. Four different outlier distributions are considered:

$$\text{M1: } \mu = (1,\ldots,1)^T, \Theta_o = \tilde{\Theta}, \quad \text{M2: } \mu = (1.5,\ldots,1.5)^T, \Theta_o = \tilde{\Theta},$$

$$\text{M3: } \mu = (1,\ldots,1)^T, \Theta_o = I_p, \quad \text{M4: } \mu = (1.5,\ldots,1.5)^T, \Theta_o = I_p.$$

We also consider the scenario where the outliers are not symmetric about the mean and simulate data from the following model

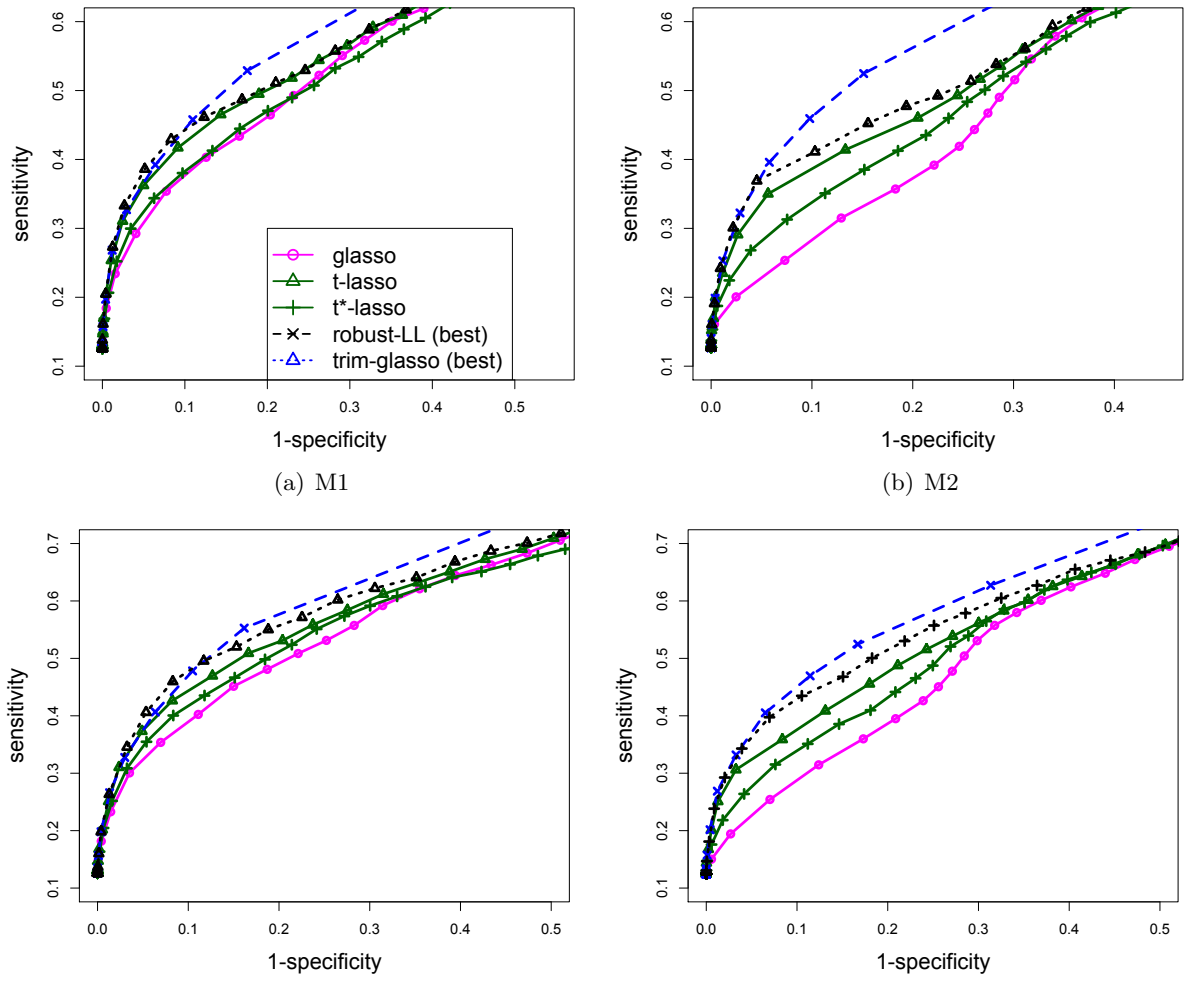

(a) M1 (b) M2

(c) M3 (d) M4

Figure 1: Average ROC curves for the comparison methods for contamination scenarios M1-M4.

M5: $y_k \sim (1-p_o)N_p(0,\Theta^{-1}) + p_o N_p(2, I_p), \quad k = 1,\ldots,n.$

For each simulation run, Θ is a randomly generated precision matrix corresponding to a network with 9 hub nodes simulated as follows. Let A be the adjacency of the network. For all $i < j$ we set $A_{ij} = 1$ with probability 0.03, and zero otherwise. We set $A_{ji} = A_{ij}$. We then randomly select 9 hub nodes and set the elements of the corresponding rows and columns of A to one with probability 0.4 and zero otherwise. Using A, the simulated nonzero coefficients of the precision matrix are sampled as follows. First we create a matrix E so that $E_{i,j} = 0$ if $A_{i,j} = 0$, and $E_{i,j}$ is sampled uniformly from $[-0.75, -0.23] \cup [0.25, 0.75]$ if $A_{i,j} \neq 0$. Then we set $E = \frac{E+E^T}{2}$. Finally we set $\Theta = E + (0.1 - \Lambda_{\min}(E))I_p$, where $\Lambda_{\min}(E)$ is the smallest eigenvalue of E. $\tilde{\Theta}$ is a randomly generated precision matrix in the same way Θ is generated.

For the robustness parameter β of the *robust-LL* method, we consider $\beta \in \{0.005, 0.01, 0.02, 0.03\}$ as recommended in [20]. For the *trim-glasso* method we consider $\frac{100h}{n} \in \{90, 85, 80\}$. Since all the

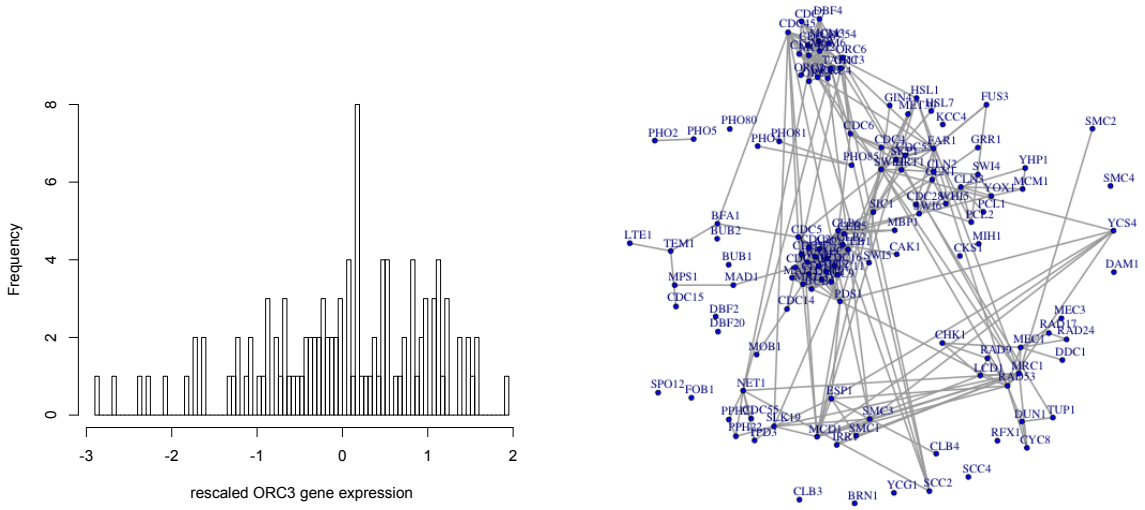

Figure 2: (a) Histogram of standardized gene expression levels for gene *ORC3*. (b) Network estimated by *trim-glasso*

robust comparison methods converge to a stationary point, we tested various initialization strategies for the concentration matrix, including I_p, $(S + \lambda I_p)^{-1}$ and the estimate from glasso. We did not observe any noticeable impact on the results.

Figure 1 presents the average ROC curves of the comparison methods over 100 simulation data sets for scenarios M1-M4 as the tuning parameter λ varies. In the figure, for *robust-LL* and *trim-glasso* methods, we depict the best curves with respect to parameter β and h respectively. Due to space constraints, the detailed results for all the values of β and h considered, as well as the results for model M5 are provided in the Supplements.

From the ROC curves we can see that our proposed approach is competitive compared the alternative robust approaches *t-lasso*, *t*-lasso* and *robust-LL*. The edge over glasso is even more pronounced for scenarios M2, M4 and M5. Surprisingly, *trim-glasso* with $h/n = 80\%$ achieves superior sensitivity for nearly any specificity.

Computationally the *trim-glasso* method is also competitive compared to alternatives. The average run-time over the path of tuning parameters λ is 45.78s for *t-lasso*, 22.14s for *t*-lasso*, 11.06s for *robust-LL*, 1.58s for trimmed lasso, 1.04s for glasso. Experiments were run on R in a single computing node with a Intel Core i5 2.5GHz CPU and 8G memory. For *t-lasso*, *t*-lasso* and *robust-LL* we used the R implementations provided by the methods' authors. For *glasso* we used the *glassopath* package.

4.2 Application to the analysis of Yeast Gene Expression Data

We analyze a yeast microarray dataset generated by [27]. The dataset concerns $n = 112$ yeast segregants (instances). We focused on $p = 126$ genes (variables) belonging to cell-cycle pathway as provided by the KEGG database [28]. For each of these genes we standardize the gene expression data to zero-mean and unit standard deviation. We observed that the expression levels of some genes are clearly not symmetric about their means and might include outliers. For example the histogram of gene *ORC3* is presented in Figure 2(a). For the *robust-LL* method we set $\beta = 0.05$

and for *trim-glasso* we use $h/n = 80\%$. We use 5-fold-CV to choose the tuning parameters for each method. After λ is chosen for each method, we rerun the methods using the full dataset to obtain the final precision matrix estimates.

Figure 2(b) shows the cell-cycle pathway estimated by our proposed method. For comparison the cell-cycle pathway from the KEGG [28] is provided in the Supplements. It is important to note that the KEGG graph corresponds to what is currently known about the pathway. It should *not* be treated as the ground truth. Certain discrepancies between KEGG and estimated graphs may also be caused by inherent limitations in the dataset used for modeling. For instance, some edges in cell-cycle pathway may not be observable from gene expression data. Additionally, the perturbation of cellular systems might not be strong enough to enable accurate inference of some of the links.

glasso tends to estimate more links than the robust methods. We postulate that the lack of robustness might result in inaccurate network reconstruction and the identification of spurious links. Robust methods tend to estimate networks that are more consistent with that from the KEGG (F_1-score of 0.23 for *glasso*, 0.37 for *t*-lasso*, 0.39 for *robust-NLL* and 0.41 for *trim-glasso*, where the F_1 score is the harmonic mean between precision and recall). For instance our approach recovers several characteristics of the KEGG pathway. For instance, genes *CDC6* (a key regulator of DNA replication playing important roles in the activation and maintenance of the checkpoint mechanisms coordinating S phase and mitosis) and *PDS1* (essential gene for meiotic progression and mitotic cell cycle arrest) are identified as a hub genes, while genes *CLB3,BRN1,YCG1* are unconnected to any other genes.

References

[1] S.L. Lauritzen. *Graphical models*. Oxford University Press, USA, 1996.

[2] Jung Hun Oh and Joseph O. Deasy. Inference of radio-responsive gene regulatory networks using the graphical lasso algorithm. *BMC Bioinformatics*, 15(S-7):S5, 2014.

[3] C. D. Manning and H. Schutze. *Foundations of Statistical Natural Language Processing*. MIT Press, 1999.

[4] J.W. Woods. Markov image modeling. *IEEE Transactions on Automatic Control*, 23:846–850, October 1978.

[5] M. Hassner and J. Sklansky. Markov random field models of digitized image texture. In *ICPR78*, pages 538–540, 1978.

[6] G. Cross and A. Jain. Markov random field texture models. *IEEE Trans. PAMI*, 5:25–39, 1983.

[7] E. Ising. Beitrag zur theorie der ferromagnetismus. *Zeitschrift für Physik*, 31:253–258, 1925.

[8] B. D. Ripley. *Spatial statistics*. Wiley, New York, 1981.

[9] E. Yang, A. C. Lozano, and P. Ravikumar. Elementary estimators for graphical models. In *Neur. Info. Proc. Sys. (NIPS)*, 27, 2014.

[10] M. Yuan and Y. Lin. Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94 (1):19–35, 2007.

[11] J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical Lasso. *Biostatistics*, 2007.

[12] O. Bannerjee, , L. El Ghaoui, and A. d'Aspremont. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Jour. Mach. Lear. Res.*, 9:485–516, March 2008.

[13] P. Ravikumar, M. J. Wainwright, G. Raskutti, and B. Yu. High-dimensional covariance estimation by minimizing ℓ_1-penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.

[14] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, UK, 2004.

[15] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the Lasso. *Annals of Statistics*, 34:1436–1462, 2006.

[16] E. Yang, P. Ravikumar, G. I. Allen, and Z. Liu. Graphical models via generalized linear models. In *Neur. Info. Proc. Sys. (NIPS)*, 25, 2012.

[17] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.

[18] Z.J. Daye, J. Chen, and Li H. High-dimensional heteroscedastic regression with an application to eqtl data analysis. *Biometrics*, 68:316–326, 2012.

[19] Michael Finegold and Mathias Drton. Robust graphical modeling of gene networks using classical and alternative t-distributions. *The Annals of Applied Statistics*, 5(2A):1057–1080, 2011.

[20] H. Sun and H. Li. Robust Gaussian graphical modeling via l1 penalization. *Biometrics*, 68:1197–206, 2012.

[21] A. Alfons, C. Croux, and S. Gelper. Sparse least trimmed squares regression for analyzing high-dimensional large data sets. *Ann. Appl. Stat.*, 7:226–248, 2013.

[22] P-L Loh and M. J. Wainwright. Regularized m-estimators with nonconvexity: Statistical and algorithmic theory for local optima. *Arxiv preprint arXiv:1305.2436v2*, 2013.

[23] C. J. Hsieh, M. Sustik, I. Dhillon, and P. Ravikumar. Sparse inverse covariance matrix estimation using quadratic approximation. In *Neur. Info. Proc. Sys. (NIPS)*, 24, 2011.

[24] Y. Nesterov. Gradient methods for minimizing composite objective function. Technical Report 76, Center for Operations Research and Econometrics (CORE), Catholic Univ. Louvain (UCL)., 2007.

[25] N. H. Nguyen and T. D. Tran. Robust Lasso with missing and grossly corrupted observations. *IEEE Trans. Info. Theory*, 59(4):2036–2058, 2013.

[26] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.

[27] Rachel B Brem and Leonid Kruglyak. The landscape of genetic complexity across 5,700 gene expression traits in yeast. *Proceedings of the National Academy of Sciences of the United States of America*, 102 (5):1572–1577, 2005.

[28] M. Kanehisa, S. Goto, Y. Sato, M. Kawashima, M. Furumichi, and M. Tanabe. Data, information, knowledge and principle: back to metabolism in kegg. *Nucleic Acids Res.*, 42:D199–D205, 2014.

[29] E. Candes, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, August 2006.

[30] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Y. Eldar and G. Kutyniok, editors, *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012. forthcoming.

Appendix

A Useful lemma(s)

Lemma 2 (Lemma 1 of [13]). *Suppose that* $\{X^{(i)}\}_{i=1}^n$ *are iid samples from* $N(0,\Sigma)$ *with* $n \geq 40\max_i \Sigma_{ii}$. *Let* $\mathcal{A}$ *be the event that*

$$\left\| \frac{1}{n}\sum_{i=1}^n X^{(i)}(X^{(i)})^\top - \Sigma \right\|_\infty \leq 8(\max_i \Sigma_{ii})\sqrt{\frac{10\tau \log p}{n}}$$

where τ *is any constant greater than 2. Then, the probability of event* $\mathcal{A}$ *occurring is at least* $1 - 4/p^{\tau-2}$.

B Proof of Theorem 1

Consider the optimal error $\widetilde{\Delta} := \widetilde{\Theta} - \Theta^*$ and $\widetilde{\Gamma} := \widetilde{w} - w^*$ where $(\widetilde{\Theta}, \widetilde{w})$ is an **arbitrary** local optimum of M-estimator (3). Our proof mainly uses the fact that $(\widetilde{\Theta}, \widetilde{w})$ is a local minimum of (3) satisfying

$$\left\langle\!\!\left\langle \frac{1}{h}\sum_{i=1}^n \widetilde{w}_i X^{(i)}(X^{(i)})^\top - (\Theta^* + \widetilde{\Delta})^{-1}, \widetilde{\Theta} - \Theta \right\rangle\!\!\right\rangle \leq -\left\langle\!\!\left\langle \partial\lambda\|\Theta^* + \widetilde{\Delta}\|_{1,\text{off}}, \widetilde{\Theta} - \Theta \right\rangle\!\!\right\rangle \quad \text{for any } \Theta \in \Omega.$$

This inequality comes from the first order stationary condition (see [22] for details) in terms of only Θ fixing w at $\widetilde{w}$. Therefore, if we take $\Theta = \Theta^*$ above, we have

$$\begin{aligned}
&\left\langle\!\!\left\langle \frac{1}{h}\sum_{i=1}^n \widetilde{w}_i X^{(i)}(X^{(i)})^\top - (\Theta^* + \widetilde{\Delta})^{-1}, \widetilde{\Theta} - \Theta^* \right\rangle\!\!\right\rangle \leq -\left\langle\!\!\left\langle \partial\lambda\|\Theta^* + \widetilde{\Delta}\|_{1,\text{off}}, \widetilde{\Theta} - \Theta^* \right\rangle\!\!\right\rangle \\
&\overset{(i)}{\leq} \lambda(\|\Theta^*\|_{1,\text{off}} - \|\widetilde{\Theta}\|_{1,\text{off}}) \leq \lambda(\|\Theta^*\|_{1,\text{off}} + \|\widetilde{\Delta}_{S^c}\|_{1,\text{off}} - \|\widetilde{\Delta}_{S^c}\|_{1,\text{off}} - \|\widetilde{\Theta}\|_{1,\text{off}}) \\
&= \lambda(\|\Theta^* + \widetilde{\Delta}_{S^c}\|_{1,\text{off}} - \|\widetilde{\Delta}_{S^c}\|_{1,\text{off}} - \|\widetilde{\Theta}\|_{1,\text{off}}) \\
&\overset{(ii)}{\leq} \lambda\big(\|\Theta^* + \widetilde{\Delta}_{S^c} + \widetilde{\Delta}_S\|_{1,\text{off}} + \|\widetilde{\Delta}_S\|_{1,\text{off}} - \|\widetilde{\Delta}_{S^c}\|_{1,\text{off}} - \|\widetilde{\Theta}\|_{1,\text{off}}\big) \\
&= \lambda(\|\widetilde{\Delta}_S\|_{1,\text{off}} - \|\widetilde{\Delta}_{S^c}\|_{1,\text{off}})\,, \qquad (9)
\end{aligned}$$

where S is true support set of Θ^*, the inequalities (i) and (ii) hold by respectively the convexity and the triangular inequality of $\|\cdot\|_{1,\text{off}}$ norm.

We first begin with the following lemma saying that $\|\widetilde{\Delta}\|_{\text{F}} \leq 1$, under the conditions of (C-3) and (C-4).

Lemma 3. *Suppose that the condition* (C-3) *and* (C-4) *hold. Moreover,* $4\max\left\{\|\frac{1}{h}\sum_{i=1}^n w_i^* X^{(i)}(X^{(i)})^\top - (\Theta^*)^{-1}\|_\infty, \tau_2(n,p)\right\} \leq \lambda \leq \frac{\kappa_l - \tau_1(n,p)}{3R}$ *as mentioned in the statement of Theorem 1. Then, for any local optimum* $(\widetilde{\Theta}, \widetilde{w})$, $\|\widetilde{\Delta}\|_{\text{F}} \leq 1$, *and the inequality* (5) *holds for all possible local optimal errors* $\widetilde{\Delta}$.

Proof. The Lemma 3 can be proved by the fact $-\log\det\Theta$ is a convex function. Hence, the function $f:[0,1]\to\mathbb{R}$ given by $f(t;\Theta^*,\widetilde{\Delta}) := -\log\det\left(\Theta^*+t\widetilde{\Delta}\right)$ is also convex in t, and $\left\langle\!\left\langle -(\Theta^*+\widetilde{\Delta})^{-1},\widetilde{\Delta}\right\rangle\!\right\rangle \geq \left\langle\!\left\langle -(\Theta^*+t\widetilde{\Delta})^{-1},\widetilde{\Delta}\right\rangle\!\right\rangle$ for $t\in[0,1]$ (see [22] for details).

Now, suppose that $\|\widetilde{\Delta}\|_{\mathrm{F}}\geq 1$. Then, we have

$$\begin{aligned}&\left\langle\!\left\langle (\Theta^*)^{-1}-(\Theta^*+\widetilde{\Delta})^{-1},\widetilde{\Delta}\right\rangle\!\right\rangle \geq \left\langle\!\left\langle (\Theta^*)^{-1}-(\Theta^*+t\widetilde{\Delta})^{-1},\widetilde{\Delta}\right\rangle\!\right\rangle\\ &=\frac{1}{t}\left\langle\!\left\langle (\Theta^*)^{-1}-(\Theta^*+t\widetilde{\Delta})^{-1},t\widetilde{\Delta}\right\rangle\!\right\rangle. \qquad (10)\end{aligned}$$

Since $\|\widetilde{\Delta}\|_{\mathrm{F}}\geq 1$, we can set $t=\frac{1}{\|\widetilde{\Delta}\|_{\mathrm{F}}}$ so that $\|t\widetilde{\Delta}\|_{\mathrm{F}}=1$. Hence, for $t\widetilde{\Delta}$, the condition (C-3) holds by assumption of the statement:

$$\left\langle\!\left\langle (\Theta^*)^{-1}-(\Theta^*+t\widetilde{\Delta})^{-1},t\widetilde{\Delta}\right\rangle\!\right\rangle \geq \kappa_l\|t\widetilde{\Delta}\|_{\mathrm{F}}^2=\kappa_l\,.$$

Combining with (10) yields

$$\left\langle\!\left\langle (\Theta^*)^{-1}-(\Theta^*+\widetilde{\Delta})^{-1},\widetilde{\Delta}\right\rangle\!\right\rangle \geq \kappa_l\|\widetilde{\Delta}\|_{\mathrm{F}}\,. \qquad (11)$$

Now, from (9) and (11) followed by the condition (C-4) and Hölder's inequity, we can obtain

$$\begin{aligned}&\kappa_l\|\widetilde{\Delta}\|_{\mathrm{F}}\leq\left\langle\!\left\langle (\Theta^*)^{-1}-\frac{1}{h}\sum_{i=1}^n\widetilde{w}_iX^{(i)}(X^{(i)})^\top,\widetilde{\Delta}\right\rangle\!\right\rangle+\lambda(\|\widetilde{\Delta}_S\|_{1,\mathrm{off}}-\|\widetilde{\Delta}_{S^c}\|_{1,\mathrm{off}})\\ &\leq\left\langle\!\left\langle (\Theta^*)^{-1}-\frac{1}{h}\sum_{i=1}^n\widetilde{w}_iX^{(i)}(X^{(i)})^\top,\widetilde{\Delta}\right\rangle\!\right\rangle+\lambda\|\widetilde{\Delta}\|_{1,\mathrm{off}}\\ &\leq\left\langle\!\left\langle (\Theta^*)^{-1}-\frac{1}{h}\sum_{i=1}^n w_i^*X^{(i)}(X^{(i)})^\top+\frac{1}{h}\sum_{i=1}^n w_i^*X^{(i)}(X^{(i)})^\top-\frac{1}{h}\sum_{i=1}^n\widetilde{w}_iX^{(i)}(X^{(i)})^\top,\widetilde{\Delta}\right\rangle\!\right\rangle+\lambda\|\widetilde{\Delta}\|_{1,\mathrm{off}}\\ &\leq\left\|\frac{1}{h}\sum_{i=1}^n w_i^*X^{(i)}(X^{(i)})^\top-(\Theta^*)^{-1}\right\|_\infty\cdot\|\widetilde{\Delta}\|_1+\tau_1(n,p)\|\widetilde{\Delta}\|_{\mathrm{F}}+\tau_2(n,p)\|\widetilde{\Delta}\|_1+\lambda\|\widetilde{\Delta}\|_{1,\mathrm{off}}\,.\end{aligned}$$

By the choice of λ in the assumption of the statement and by the fact that $\|\widetilde{\Delta}\|_{1,\mathrm{off}}\leq\|\widetilde{\Delta}\|_1$ and $\|\widetilde{\Delta}\|_1\leq\|\widetilde{\Theta}\|_1+\|\Theta^*\|_1\leq 2R$, we can rearrange the above inequality into

$$\|\widetilde{\Delta}\|_{\mathrm{F}}\leq\frac{3\lambda}{2\big(\kappa_l-\tau_1(n,p)\big)}\|\widetilde{\Delta}\|_1\leq\frac{3\lambda R}{\big(\kappa_l-\tau_1(n,p)\big)}\leq 1\,,$$

which conflicts with the assumption in the beginning of this proof. Hence, by contradiction, we can conclude $\|\widetilde{\Delta}\|_{\mathrm{F}}\leq 1$ under conditions in the statement. □

Now, by the RSC condition in (C-3) for $\widetilde{\Delta}$, we obtain

$$\begin{aligned}&\kappa_l\|\widetilde{\Delta}\|_{\mathrm{F}}^2\leq\left\langle\!\left\langle (\Theta^*)^{-1}-(\Theta^*+\widetilde{\Delta})^{-1},\widetilde{\Delta}\right\rangle\!\right\rangle\\ &=\Big\langle\!\Big\langle\underbrace{\frac{1}{h}\sum_{i=1}^n w_i^*X^{(i)}(X^{(i)})^\top}_{\text{(I)}}-\underbrace{\frac{1}{h}\sum_{i=1}^n w_i^*X^{(i)}(X^{(i)})^\top}_{\text{(II)}}+\underbrace{\frac{1}{h}\sum_{i=1}^n\widetilde{w}_iX^{(i)}(X^{(i)})^\top}_{\text{(III)}}-\underbrace{\frac{1}{h}\sum_{i=1}^n\widetilde{w}_iX^{(i)}(X^{(i)})^\top}_{\text{(I)}}\\ &\quad+\underbrace{(\Theta^*)^{-1}}_{\text{(II)}}-\underbrace{(\Theta^*+\widetilde{\Delta})^{-1}}_{\text{(III)}},\widetilde{\Delta}\Big\rangle\!\Big\rangle\,.\end{aligned}$$

By the condition (C-4)

$$\text{(I): } \Big\langle\!\Big\langle \frac{1}{h}\sum_{i=1}^{n} w_i^* X^{(i)}(X^{(i)})^\top - \frac{1}{h}\sum_{i=1}^{n} \widetilde{w}_i X^{(i)}(X^{(i)})^\top,\, \widetilde{\Delta} \Big\rangle\!\Big\rangle = -\Big\langle\!\Big\langle \frac{1}{h}\sum_{i=1}^{n} \widetilde{\Gamma}_i X^{(i)}(X^{(i)})^\top,\, \widetilde{\Delta} \Big\rangle\!\Big\rangle$$
$$\leq \tau_1(n,p)\|\widetilde{\Delta}\|_{\mathrm{F}} + \tau_2(n,p)\|\widetilde{\Delta}\|_1 \leq \tau_1(n,p)\|\widetilde{\Delta}\|_{\mathrm{F}} + \frac{\lambda}{4}\|\widetilde{\Delta}\|_1\,,$$

and at the same time, by (9),

$$\text{(III): } \Big\langle\!\Big\langle \frac{1}{h}\sum_{i=1}^{n} \widetilde{w}_i X^{(i)}(X^{(i)})^\top - \big(\Theta^* + \widetilde{\Delta}\big)^{-1},\, \widetilde{\Delta} \Big\rangle\!\Big\rangle \leq \lambda(\|\widetilde{\Delta}_S\|_{1,\mathrm{off}} - \|\widetilde{\Delta}_{S^c}\|_{1,\mathrm{off}})\,.$$

Finally, by the assumption on λ and Hölder's inequity,

$$\text{(II): } \Big\langle\!\Big\langle -\frac{1}{h}\sum_{i=1}^{n} w_i^* X^{(i)}(X^{(i)})^\top + \big(\Theta^*\big)^{-1},\, \widetilde{\Delta} \Big\rangle\!\Big\rangle$$
$$\leq \Big\| \frac{1}{h}\sum_{i=1}^{n} w_i^* X^{(i)}(X^{(i)})^\top - (\Theta^*)^{-1} \Big\|_\infty \cdot \|\widetilde{\Delta}\|_1 \leq \frac{\lambda}{4}\|\widetilde{\Delta}\|_1\,.$$

Combining all pieces together yields

$$0 \leq \kappa_l \|\widetilde{\Delta}\|_{\mathrm{F}}^2 \leq \frac{3\lambda}{2}\|\widetilde{\Delta}_S\|_1 + \frac{\lambda}{2}\|\widetilde{\Delta}_{S^c}\|_1 - \lambda\|\widetilde{\Delta}_{S^c}\|_{1,\mathrm{off}} + \tau_1(n,p)\|\widetilde{\Delta}\|_{\mathrm{F}}$$
$$\leq \frac{3\lambda}{2}\|\widetilde{\Delta}_S\|_1 + \frac{\lambda}{2}\|\widetilde{\Delta}_{S^c}\|_{1,\mathrm{diag}} - \frac{\lambda}{2}\|\widetilde{\Delta}_{S^c}\|_{1,\mathrm{off}} + \tau_1(n,p)\|\widetilde{\Delta}\|_{\mathrm{F}} \tag{12}$$

where $\|\cdot\|_{1,\mathrm{diag}}$ is defined as the absolute sum of diagonal elements of a matrix, and $\|\widetilde{\Delta}_{S^c}\|_{1,\mathrm{diag}} \leq \sqrt{p}\|\widetilde{\Delta}_{S^c}\|_{\mathrm{F}} \leq \sqrt{p}\|\widetilde{\Delta}\|_{\mathrm{F}}$. Therefore, (12) implies

$$\|\widetilde{\Delta}\|_{\mathrm{F}} \leq \frac{1}{\kappa_l}\Big(2\lambda\sqrt{k+p} + \tau_1(n,p)\Big)\,.$$

At the same time in order to derive $\|\cdot\|_{1,\mathrm{off}}$ error bound, we again use the inequality in (12):

$$\|\widetilde{\Delta}_{S^c}\|_{1,\mathrm{off}} \leq 3\|\widetilde{\Delta}_S\|_{1,\mathrm{off}} + \|\widetilde{\Delta}_{S^c}\|_{1,\mathrm{diag}} + \frac{2}{\lambda}\tau_1(n,p)\|\widetilde{\Delta}\|_{\mathrm{F}}\,.$$

Hence,

$$\|\widetilde{\Delta}\|_{1,\mathrm{off}}$$
$$\leq \|\widetilde{\Delta}_S\|_{1,\mathrm{off}} + \|\widetilde{\Delta}_{S^c}\|_{1,\mathrm{off}} \leq 4\|\widetilde{\Delta}_S\|_{1,\mathrm{off}} + \|\widetilde{\Delta}_{S^c}\|_{1,\mathrm{diag}} + \frac{2}{\lambda}\tau_1(n,p)\|\widetilde{\Delta}\|_{\mathrm{F}}$$
$$\leq 5\sqrt{k+p}\|\widetilde{\Delta}\|_{\mathrm{F}} + \frac{2}{\lambda}\tau_1(n,p)\|\widetilde{\Delta}\|_{\mathrm{F}} = \Big(5\sqrt{k+p} + \frac{2}{\lambda}\tau_1(n,p)\Big)\|\widetilde{\Delta}\|_{\mathrm{F}}$$
$$\leq \frac{2}{\lambda\,\kappa_l}\Big(3\lambda\sqrt{k+p} + \tau_1(n,p)\Big)^2\,,$$

which completes the proof.

C Proof of Corollary 1

In order to utilize Theorem 1, we simply need to specify the quantity $\tau_1(n,p)$ and $\tau_2(n,p)$ in (C-4). Toward this, we follow similar strategy as in [25]: given Δ, we divide the index of Δ into the disjoint exhaustive subsets $S_1, S_2, \ldots, S_q$ of size $|B|$ such that S_1 contains $|B|$ largest absolute elements in Δ, and so on. Then, we have

$$\begin{aligned}
&\Big|\Big\langle\!\Big\langle \sum_{i=1}^{n} \widetilde{\Gamma}_i X^{(i)}(X^{(i)})^\top,\, \widetilde{\Delta} \Big\rangle\!\Big\rangle\Big| = \Big|\Big\langle\!\Big\langle \sum_{i\in B} \widetilde{\Gamma}_i X^{(i)}(X^{(i)})^\top,\, \widetilde{\Delta} \Big\rangle\!\Big\rangle\Big| \\
&= \Big|\sum_{j=1}^{q} \Big\langle\!\Big\langle \sum_{i\in B} \widetilde{\Gamma}_i \big[X^{(i)}(X^{(i)})^\top\big]_{S_j},\, [\widetilde{\Delta}]_{S_j} \Big\rangle\!\Big\rangle\Big| \le \sum_{j=1}^{q} \Big|\Big\langle\!\Big\langle \sum_{i\in B} \widetilde{\Gamma}_i \big[X^{(i)}(X^{(i)})^\top\big]_{S_j},\, [\widetilde{\Delta}]_{S_j} \Big\rangle\!\Big\rangle\Big| .
\end{aligned}$$

Let $D_{\widetilde{\Gamma}}$ be a $|B| \times |B|$ diagonal matrix whose i-th diagonal entry is $[D_{\widetilde{\Gamma}}]_{ii} := \widetilde{\Gamma}_i$. Let also X^B is a $|B| \times p$ design matrix for samples in the set B. Finally $X^B_{S_j}$ denotes a $|B| \times |S_j|$ sub-matrix of X^B whose columns are indexed by S_j. Then,

$$\begin{aligned}
&\sum_{j=1}^{q} \Big|\Big\langle\!\Big\langle \sum_{i\in B} \widetilde{\Gamma}_i \big[X^{(i)}(X^{(i)})^\top\big]_{S_j},\, [\widetilde{\Delta}]_{S_j} \Big\rangle\!\Big\rangle\Big| = \sum_{j=1}^{q} \Big|\mathrm{Trace}\Big([\widetilde{\Delta}]_{S_j}^\top [X^B_{S_j}]^\top D_{\widetilde{\Gamma}} X^B_{S_j}\Big)\Big| \\
&= \sum_{j=1}^{q} \Big|\Big\langle\!\Big\langle X^B_{S_j}[\widetilde{\Delta}]_{S_j},\, D_{\widetilde{\Gamma}} X^B_{S_j} \Big\rangle\!\Big\rangle\Big| \le \sum_{j=1}^{q} \|X^B_{S_j}[\widetilde{\Delta}]_{S_j}\|_{\mathrm{F}}\, \|D_{\widetilde{\Gamma}} X^B_{S_j}\|_{\mathrm{F}} \\
&\le \sqrt{|B|}\Big(\max_j |\!|\!| X^B_{S_j} |\!|\!|_2\Big)^2 \underbrace{\Big(\textstyle\sum_j \|[\widetilde{\Delta}]_{S_j}\|_{\mathrm{F}}\Big)}_{\text{(I)}} .
\end{aligned}$$

(I): by the standard bound in [29], we obtain

$$\sum_j \|[\widetilde{\Delta}]_{S_j}\|_{\mathrm{F}} = \|[\widetilde{\Delta}]_{S_1}\|_{\mathrm{F}} + \sum_{j=2}^{q} \|[\widetilde{\Delta}]_{S_j}\|_1 \le \|[\widetilde{\Delta}]_{S_1}\|_{\mathrm{F}} + \frac{1}{\sqrt{|B|}} \sum_{j=2}^{q} \|[\widetilde{\Delta}]_{S_j}\|_1 \le \|\widetilde{\Delta}\|_{\mathrm{F}} + \frac{1}{\sqrt{|B|}} \|\widetilde{\Delta}\|_1 .$$

Combining all pieces together yields

$$\Big|\Big\langle\!\Big\langle \frac{1}{h}\sum_{i=1}^{n} \widetilde{\Gamma}_i X^{(i)}(X^{(i)})^\top,\, \widetilde{\Delta} \Big\rangle\!\Big\rangle\Big| \le \frac{\sqrt{|B|}}{h}\Big(\max_j |\!|\!| X^B_{S_j} |\!|\!|_2\Big)^2 \Big(\|\widetilde{\Delta}\|_{\mathrm{F}} + \frac{1}{\sqrt{|B|}}\|\widetilde{\Delta}\|_1\Big),$$

hence, we can guarantee the condition (C-4) with functions

$$\begin{aligned}
\tau_1(n,p) &= f(X^B)\sqrt{\frac{|B|\log p}{h}} \quad \text{and} \\
\tau_2(n,p) &= f(X^B)\sqrt{\frac{\log p}{h}} .
\end{aligned}$$

To complete the proof, we need to specify the quantity $\big\|\frac{1}{h}\sum_{i=1}^{n} w_i^* X^{(i)}(X^{(i)})^\top - (\Theta^*)^{-1}\big\|_\infty = \big\|\frac{1}{h}\sum_{i\in G} w_i^* X^{(i)}(X^{(i)})^\top - (\Theta^*)^{-1}\big\|_\infty$ for the appropriate choice of λ as stated in Theorem 1. Recall that we constructed w^* as follows: w_i^* is simply set to $\widetilde{w}_i$ if $i \in G$, and $w_i^* = 0$ for $i \in B$. Let

$\bar{G}$ be the subset of G such that $w_i^* = 1$ and $\bar{G}^c$ be the subset such that $w_i^* = 0$. Then, we have $h \geq |\bar{G}| \geq h - |B|$, and hence $h - |\bar{G}| \leq |B|$. Now, by Lemma 2, we can obtain the following bound:

$$\begin{aligned}
&\Big\|\frac{1}{h}\sum_{i=1}^{n} w_i^* X^{(i)}(X^{(i)})^\top - (\Theta^*)^{-1}\Big\|_\infty = \Big\|\frac{|\bar{G}|}{h}\frac{1}{|\bar{G}|}\sum_{i\in\bar{G}} X^{(i)}(X^{(i)})^\top - (\Theta^*)^{-1}\Big\|_\infty \\
&= \Big\|\frac{|\bar{G}|}{h}\Big(\frac{1}{|\bar{G}|}\sum_{i\in\bar{G}} X^{(i)}(X^{(i)})^\top - (\Theta^*)^{-1}\Big) - \Big(\frac{h-|\bar{G}|}{h}\Big)(\Theta^*)^{-1}\Big\|_\infty \\
&\leq \Big\|\frac{|\bar{G}|}{h}\Big(\frac{1}{|\bar{G}|}\sum_{i\in\bar{G}} X^{(i)}(X^{(i)})^\top - (\Theta^*)^{-1}\Big)\Big\|_\infty + \Big\|\Big(\frac{h-|\bar{G}|}{h}\Big)(\Theta^*)^{-1}\Big\|_\infty \\
&\leq \Big\|\frac{1}{|\bar{G}|}\sum_{i\in\bar{G}} X^{(i)}(X^{(i)})^\top - (\Theta^*)^{-1}\Big\|_\infty + \frac{|B|}{h}\|\Sigma^*\|_\infty \\
&\leq 8(\max_i \Sigma^*_{ii})\sqrt{\frac{10\tau\log p}{\bar{G}}} + \frac{|B|}{h}\|\Sigma^*\|_\infty \leq 8(\max_i \Sigma^*_{ii})\sqrt{\frac{10\tau\log p}{h-|B|}} + \frac{|B|}{h}\|\Sigma^*\|_\infty
\end{aligned}$$

with probability at least $1 - 4/p^{\tau-2}$ for any $\tau > 2$.

D Proof of Corollary 2

Under (C-B2), $h - |B| \geq (n - a\sqrt{n}) - a\sqrt{n}$. Hence, if $n \geq 16a^2$, then $h - |B| \geq (n - a\sqrt{n}) - a\sqrt{n} \geq \frac{n}{2}$. Moreover, $\frac{|B|}{h} \leq \frac{a\sqrt{n}}{n/2} \leq \frac{2a}{\sqrt{n}}$. Therefore, from the Corollary 1, the selection of λ in the statement satisfies $\lambda \geq 4\max\big\{\|\frac{1}{h}\sum_{i=1}^n w_i^* X^{(i)}(X^{(i)})^\top - (\Theta^*)^{-1}\|_\infty\,,\, \tau_2(n,p)\big\}$.

Furthermore, as long as $n \geq \big(|||\Theta^*|||_2 + 1\big)^4\big(3Rc + f(X^B)\sqrt{2|B|}\big)^2(\log p)$,

$$\lambda = c\sqrt{\frac{\log p}{n}} \leq \frac{\big(|||\Theta^*|||_2 + 1\big)^{-2} - f(X^B)\sqrt{\frac{2|B|\log p}{n}}}{3R}\,,$$

and therefore we have $\lambda \leq \frac{\kappa_l - f(X^B)\sqrt{\frac{|B|\log p}{h}}}{3R}$ where c is defined as $4\max\big\{16(\max_i \Sigma^*_{ii})\sqrt{5\tau} + \frac{2a\|\Sigma^*\|_\infty}{\sqrt{\log p}}\,,\, \sqrt{2}f(X^B)\big\}$, as stated.

E Proof of Corollary 3

In this proof, we simply need to specify the quantity $f(X^B)$ under the condition (C-B2) and (C-B3), and then we can appeal to the result in Corollary 2.

Provided $|B| \geq \exp(1)$, $\binom{p}{|B|} \leq \big(\frac{\exp(1)p}{|B|}\big)^{|B|} \leq p^{|B|}$. As discussed in [30, 25], if (C-B3) holds, for every $t > 0$, we have

$$\frac{1}{\sqrt{|B|}}\max_j |||X^B_{S_j}|||_2 \leq \sqrt{|||\Sigma_B|||_2}\,(2 + t)$$

with probability at least $1 - 2\binom{p}{|B|}\exp(-t^2|B|/2) \geq 1 - 2\exp\big(-t^2|B|/2 + |B|\log p\big)$. Setting $t = 2\sqrt{\log p}$, we obtain

$$\max_j \|X^B_{S_j}\|_2 \leq 2(1+\sqrt{\log p})\sqrt{\|\Sigma_B\|_2}\sqrt{|B|}$$

with probability $1 - p^{-|B|}$. Therefore, under (C-B2),

$$\Big(\max_j \|X^B_{S_j}\|_2\Big)^2 \leq 4\big(1+\sqrt{\log p}\big)^2\|\Sigma_B\|_2|B| \leq 4a\big(1+\sqrt{\log p}\big)^2\|\Sigma_B\|_2\sqrt{n}$$
$$= \frac{4a\big(1+\sqrt{\log p}\big)^2\|\Sigma_B\|_2\sqrt{n}}{\sqrt{h\log p}}\sqrt{h\log p} \leq \frac{4\sqrt{2}a\big(1+\sqrt{\log p}\big)^2\|\Sigma_B\|_2}{\sqrt{\log p}}\sqrt{h\log p},$$

as specified in the statement.

F Additional Results

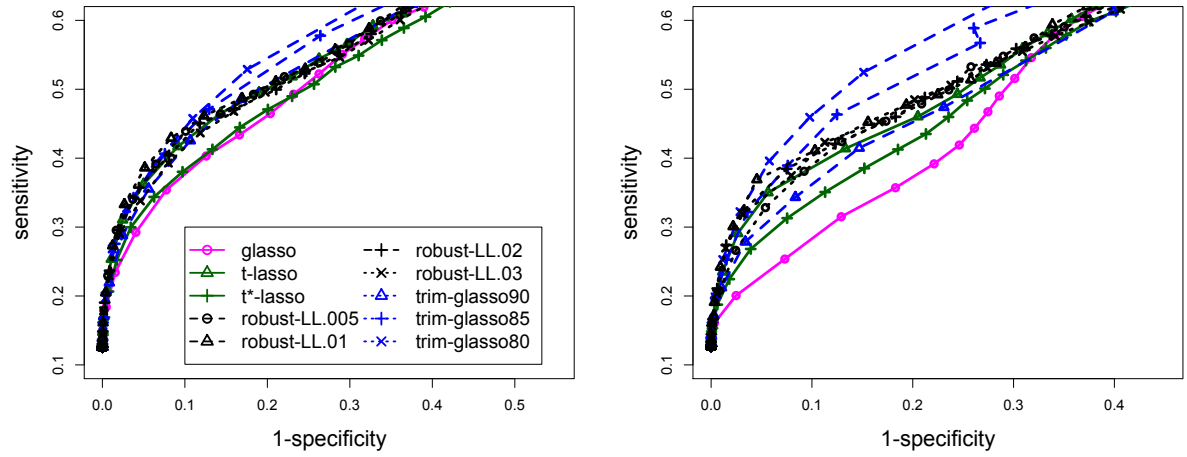

(a) M1

(b) M2

(c) M3

(d) M4

Figure 3: Average ROC curves for the comparison methods for contamination scenarios M1-M4.

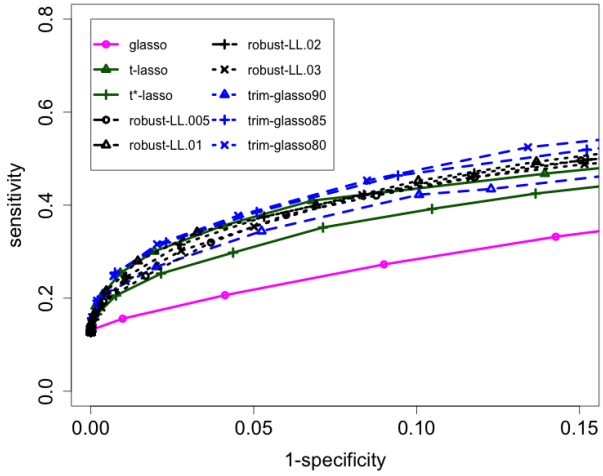

Figure 4: Average ROC curves for the comparison methods for contamination scenario M5.

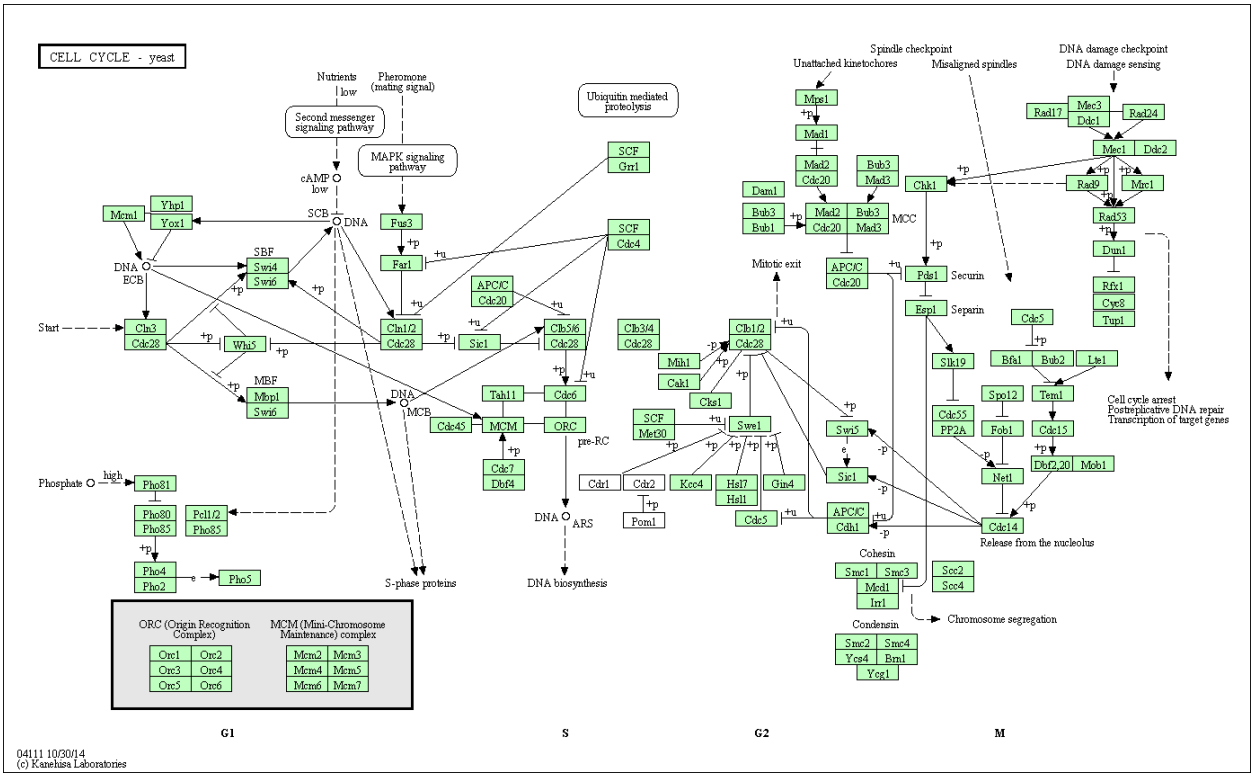

Figure 5: Reference Yeast Cell Signaling Network from the KEGG database [28].